

Remarks

Entry of the amendments, reconsideration of the application, as amended, and allowance of all pending claims are respectfully requested. Claims 1-49 remain pending. **Applicants respectfully request that these Remarks be carefully considered by the Examiner.**

With the above amendments, applicants are removing the language objected to by the Examiner. This amendment should be entered, even though it is after final, since the language being removed is just a further clarification of the language that is already in the claim. In particular, the claim indicates that there is an individual state and a group state. The individual state is individual to the individual prospective member and the group state is the state defined for the processing group. The language being deleted is language that applicants thought made the individual state more clear, but since it has caused some confusion, applicants are deleting this language. However, the claim still recites an individual state that defines state for the individual member and a group state that defines state for the processing group. Thus, applicants respectfully request that this amendment be entered.

In the Office Action dated March 16, 2004, claims 1, 3, 5, 28, 35 and 42 are rejected under 35 U.S.C. 112, first paragraph, as failing to comply with the written description requirement. Although applicants strenuously disagree with this rejection, applicants have deleted the rejected language in a bona fide attempt to further prosecution of this application. Applicants specification clearly indicates that there are two states, an individual state and a group state. It is inherent in the word individual that the individual state is the state for a particular individual and further that the group state is the state defined for the group. Applicants compare state of the individual with state of the group and perform processing based on the results of the comparison.

Based on the amendment and/or the remarks, applicants respectfully request withdrawal of the §112 rejection.

In addition to the above, the pending claims have been rejected. However, before applicants address the particular rejections, applications would like to address several remarks made in the Response to Argument section. It is stated in that section that applicants arguments

fail to comply with 37 CFR 1.111(b) and (c) because they do not specifically point out the language of the claims that patentably distinguishes them from the references. Applicants strenuously disagree with this statement. In the previous Response to Office Action, applicants specifically indicated the language that they believed distinguished their invention from the cited art. Also, applicants specifically indicated the teachings in the reference that taught other than what is claimed by applicants. Therefore, applicants have clearly pointed out the language that distinguishes their invention. Applicants will again point out the novel features below.

Claims 1-21 and 28-49 are rejected under 35 U.S.C. 102(c) as being anticipated by Mojin (U.S. Patent No. 6,108,699). Applicants respectfully, but most strenuously, traverse this rejection for the reasons herein.

Applicants' invention is directed, in one aspect, to a protocol used to join a prospective member to a processing group. The join protocol includes various steps taken to ensure configuration consistency. These steps include, for instance, the comparison and updating of state to appropriately join the prospective member to the group. One example of the state used in applicants' protocol is individual prospective member state (i.e., state specific to the prospective member). In one instance, this state is a sequence number of the prospective member that is used throughout the join protocol to control whether a prospective member joins the group and to ensure configuration consistency.

As one particular example, applicants claim a method of managing processing groups of a distributed computing environment (e.g., independent claim 1). The method includes, for instance, comparing at least a portion of an individual prospective member state of a prospective member of a processing group with at least a portion of a group state of the processing group, the individual prospective member state comprising state defined for the individual prospective member; updating the at least a portion of the individual prospective member state, should the comparing indicate a difference; and joining the prospective member to the processing group, in response to the updating. Thus, in applicants' claimed invention, the individual prospective member state is employed in the joining. It is the individual prospective member state that is updated when the comparing indicates a difference. The individual prospective member state

helps ensure the consistency of the configuration. The use of the individual prospective member state in the joining process is very different from the teachings of Moiin.

While Moiin describes a join protocol, the join protocol of Moiin is distinct from the protocol of applicants' claimed invention. For instance, the join protocol of Moiin fails to employ an individual prospective member state, as claimed by applicants. Instead, Moiin only uses group state. That is, the comparisons and updating in Moiin are of group state. There is no teaching or suggestion of comparing and/or updating individual member state, as claimed by applicants.

To further explain, FIG. 4 and Cols. 5-7 of Moiin describe that the join process of Moiin includes broadcasting reconfiguration messages to nodes 1-5; waiting to receive replies to those reconfiguration messages; and eventually updating various group states, including a next cluster size field, a next cluster vector, a cluster size field and a cluster vector. Each of the states that is compared and updated defines the group. For instance, the cluster vector field describes all of the nodes of the cluster, and the cluster size field describes the size of the cluster. None of those fields includes state defined for the individual member, as claimed by applicants. The only individual state described in Moiin is the identification field, and this field is not compared or updated.

The comparing and updating of group state is described throughout Moiin. For instance, it is explicitly stated in Col. 6, lines 13-17:

In step 408, CMM 220A (FIG. 3) updates next cluster size field 308 and next cluster vector 310 to represent a cluster which includes node 0 and all nodes from which CMM 220A receives a reconfiguration message in step 406 (FIG. 4).

It further states in Col. 6, lines 35-42:

Specifically, in test step 410 (FIG. 4), CMM 220A (FIG. 3) compares the cluster size represented in cluster size field 304 to a value of one to determine whether any node other than node 0 is a member of the prospective cluster. If the cluster size is greater than one, processing transfers to step 414 (FIG. 4) which is described below. Conversely, if the cluster size is not greater than one, processing transfers to test step 412.

Yet further, Col. 7, lines 44-60 state:

In test step 422, CMM 220A (FIG. 3) compares the received reconfiguration messages to determine whether all the received reconfiguration messages represent exactly the same cluster, i.e., whether all received reconfiguration messages agree as to cluster membership in the prospective cluster. If any of the received reconfiguration messages do not agree as to cluster membership, processing transfers from test step 422 (FIG. 4) to step 420 in which the reconfiguration of the cluster fails in the manner described above. Conversely, if all received reconfiguration messages agree as to membership in the prospective, processing transfers from test step 422 to step 424. In step 424, the prospective cluster is accepted and node 0 saves the prospective cluster as the current cluster by copying data from next cluster size field 308 (FIG. 3) and next cluster vector field 310 to cluster size field 304 and cluster vector field 306, respectively.

Thus, it is repeatedly stated in Moiin that the comparisons and updates are of group state, i.e., state that defines the group, rather than of individualized state of the prospective member. Again, the only individualized state described in Moiin is identification field 302 (FIG. 3), which is described in Col. 5, lines 33-35. Moiin states: "This identification field includes data which uniquely identifies node 0 and distinguishes node 0 from nodes 1-5 (FIG. 1)." Thus, while Moiin describes an identifier field, there is no description, teaching or suggestion of a comparison of the identifier field with the group state. Further, there is certainly no teaching or suggestion of updating the identifier field. Therefore, there is no description, teaching or suggestion in Moiin of comparing at least a portion of an individual prospective member state of a prospective member with the at least a portion of the group state, in which the individual prospective member state comprises state defined for the individual prospective member; nor is there any description, teaching or suggestion of updating the individual prospective member state. Thus, applicants respectfully submit that Moiin does not anticipate independent claim 1, as well as independent claims 3 and 5. Further, for similar reasons, applicants respectfully submit that Moiin does not anticipate independent claims 28, 35 and 42.

The dependent claims are patentable for the same reasons as the independent claims, as well as for their own additional features. For example, dependent claim 7 specifically indicates that the individual prospective member state comprises a sequence number of the prospective

member, and the comparing compares the sequence number of the prospective member with a sequence number of the processing group. Applicants respectfully submit that there is no teaching or suggestion in Moiin of a sequence number. As understood, a sequence is defined, for instance, as a continuity of progression (see, e.g., Webster's Ninth New Collegiate Dictionary). Applicants use the sequence number to identify, for instance, the particular version of the group configuration. Applicants respectfully submit that there is no teaching or suggestion of a sequence number in Moiin.

Support for the rejection of claim 7 is indicated at Col. 5, lines 24-26 and 32-46 and Col. 6, lines 13-17. However, a careful reading of those sections fails to describe a sequence number. Instead, various fields are described without reference to sequence. For instance, an identification field is described which uniquely describes the node. This identification field, however, is not described as having a particular sequence or of being a sequence number. Further, this field is neither compared nor updated. Additionally, a cluster vector field and a cluster vector size are described. These fields describe the members of the cluster and the size of the cluster. They do not describe a sequence number of the prospective member or of the group. Based on the foregoing, applicants respectfully submit that dependent claim 7, and other similarly rejected claims, are not anticipated by Moiin.

Further, dependent claim 49 specifically states that the sequence number of the prospective member identifies a version of a proposed processing group to join. There is no description, teaching or suggestion in Moiin of a sequence number identifying a version of a processing group. Thus, claim 49 is not anticipated, taught or suggested by Moiin, and applicants respectfully request an indication of allowability for dependent claim 49.

In addition to the above, claims 22-27 are rejected under 35 U.S.C. 102(e) as being anticipated by Shrivastava et al. (U.S. Patent No. 6,449,734). Applicants respectfully, but most strenuously, traverse this rejection.

This aspect of applicants' invention is directed to handling failed members. As one example, applicants claim a method of managing processing groups of a distributed computing environment (e.g., in independent claim 22), in which the method includes detecting a failure of at least one member of a processing group; quiescing activity to a group state of the processing

group; and updating at least a portion of the group state in order to exclude the at least one member of the processing group. The updating includes updating a sequence number of the group state, the sequence number identifying a version of the processing group. Thus, in this aspect of applicants' claimed invention, when a failure of a member of a processing group is detected, the processing group is updated in order to exclude that member. The updating includes updating the sequence number that identifies a version of the processing group. This is very different from the teachings of Shrivastava.

Although Shrivastava uses similar words, such as sequence number or failure, Shrivastava still does not describe, teach or suggest one or more aspects of applicants' claimed invention. In applicants' claimed invention, a sequence number identifies a version of the processing group (e.g., cluster) and the sequence number is updated in order to exclude a member of the processing group. In contrast, in Shrivastava, a sequence number is used to identify a particular transaction (see, e.g., Col. 13, lines 37-45) and not to describe the version of a processing group. Thus, the sequence number of Shrivastava is very different from the sequence number of applicants' claimed invention.

Further, the sequence number in Shrivastava is not updated when a transaction is not added to the group of transactions. Instead, the same transaction number is used and it is the duplicate sequence number that indicates that a transaction failed (see, e.g., Abstract; Col. 123, lines 53-49). The failure or exclusion is not indicated by an update of the group state, as claimed by applicants. There is no update in Shrivastava. Instead, the same, non-updated number is used.

Moreover, although various references are made in Shrivastava of removing a system from a cluster (see, e.g., Col. 5, lines 53-57; Col. 9, lines 32-33), Shrivastava does not describe how to remove that system from the cluster. Instead, it focuses on how to keep the transaction consistent. There is absolutely no teaching or suggestion in Shrivastava of detecting a failure of a member of a processing group and updating group state in order to exclude that member from the processing group, wherein the updating comprises updating a sequence number of the group, in which the sequence number identifies a version of the group. Since this is missing from Shrivastava, Shrivastava fails to describe, teach or suggest applicants' claimed invention.

For all of the above reasons, applicants respectfully submit that Mojin and Shrivastava fail to describe, teach or suggest applicants' claimed invention, and applicants respectfully request an indication of allowance for all pending claims.

Should the Examiner wish to discuss this case with applicants' attorney, please contact applicants' attorney at the below listed number.

Respectfully submitted,

Blanche E. Schiller
Blanche E. Schiller
Attorney for Applicants
Registration No.: 35,670

Dated: May 17, 2004.

HESLIN ROTHENBERG FARLEY & MESITI P.C.
5 Columbia Circle
Albany, New York 12203-5160
Telephone: (518) 452-5600
Facsimile: (518) 452-5579